

BOG 10: OS/R Node Local

ASCR Workshop on Extreme Heterogeneity in HPC
23-25 Jan 2018

Moderator(s): Michael Lang

BOG X Contributors

BOGists:

- Barry Rountree rountree@llnl.gov
- Beard, Jonathan jonathan.beard@arm.com
- Brian Kocoloski brian.kocoloski@wustl.edu
- Chris Rossbach rossbach@cs.utexas.edu
- Christian Engelmann engelmann@computer.org
- David Bernholdt bernholdtde@ornl.gov
- Gokhale, Maya B. gokhale2@llnl.gov
- Kenneth Moreland kmorel@sandia.gov
- Lang, Mike mlang@lanl.gov
- Larry Kaplan lkaplan@cray.com
- Lofstead, Gerald F II gflfst@sandia.gov
- Mattan mattan.erez@utexas.edu
- Noah Watkins jayhawk@cs.ucsc.edu
- Nowell, Lucy Lucy.Nowell@science.doe.gov
- Paul Peltz Jr peltz@lanl.gov
- Ron Brightwell rbbrigh@sandia.gov
- sriram krishnamoorthy sriram.krishnamoorthy@gmail.com
- Suren Byna sbyna@lbl.gov
- Sweeney, Christine cahrens@lanl.gov
- Terry Jones trjones@ornl.gov
- Vetter, Jeffrey S. vetter@computer.org
- Yoonho Park yooho@us.ibm.com
- Doug Jacobsen dmjacobsen@lbl.gov
- Tina Declerck tmdeclerck@lbl.gov
- Rebecca Hartman-Baker rjhartmanbaker@lbl.gov
- Latchesar Ionkov lionkov@lanl.gov
- Vitus Leung
- Pete Beckman\
- Kamil Iskra
- Andrew J. Younge
- Brian Van Essen
- Brian G. G. G.

BOG 10 OS/R Challenge Assessment

Discussion to identify research required to get from where the capabilities are now to where they need to be by 2030.

Memory - storage is blurring

Definition of a Node/Rack is blurring

Many distributed processing elements of differing characteristics (power, sequential, parallel, specialized, reconfigurable, BW, Latency...)

Mostly Static Resources vs Dynamic (FPGA can have 2 cores, then 10 cores)

Trapped between two ill defined areas - Hardware and Programming models

BOG 10 OS/R Targets for 2030 (**PRDs)

1. OS-R / Hardware co-design for unified heterogeneous components: computational, active memory, network, and storage components
 - a. Virtual memory system redesign
 - b. Scheduling/discovery/composition/arbitration across dynamic heterogeneous components
 - c. Non-Linux OS alternatives, non POSIX, rearchitecting control plane / data plane boundaries
 - d. Virtualization and containerization
2. Decentralized Resource Management
 - a. Interfaces between multiple local OSes; interfaces between LOS and GOS
3. Applying ML/AI to OS and runtime resource management
 - a. Leveraging accelerators for OS/R needs including global OS
 - b. autonomy

BOG 10 Capability Targets for Extreme Heterogeneity

BOG X brainstorming and discussion of capabilities that will be needed in the 2025-2035 timeframe to make increasingly heterogeneous hardware technologies useful and productive for science applications.

BOG 10 OS/R

What does an OS do -- local resource manager:

Processing, Memory, Communication, Storage

What is changing:

Processing, Memory, Communication, Storage

BOG 10 OS/R Targets for 2030

1. OS / Hardware co-design for unified heterogeneous components: computational, active memory, network, and storage components
 - a. Virtual memory system redesign
 - b. Scheduling/discovery/composition/arbitration across dynamic heterogeneous components
2. Non-Linux OS alternatives, non POSIX, rearchitecting control plane / data plane boundaries
3. Interfaces between multiple local OSES; interfaces between LOS and GOS
4. Virtualization and containerization
5. Applying ML/AI to OS and runtime resource management
 - a. Leveraging accelerators for OS needs

BOG 10 OS/R Current Status

BOG X survey describing the current status of science, technology, or practice related to this theme (starting with material directly from the FSD).

BOG 10 OS/R Current Capability

Capability 1: Basic support for accelerators

- Fine-grained support of CPU, memory

- Coarse-grained support of Accelerators

Capability 2: limited interaction/communication between LOS and GOS

BOG 10 OS/R Challenge Assessment

Discussion to identify research required to get from where the capabilities are now to where to where they need to be by 2030.

Memory - storage is blurring

Definition of a Node/Rack is blurring

Many distributed processing elements of differing characteristics (power, sequential, parallel, specialized, reconfigurable, BW, Latency...)

Mostly Static Resources vs Dynamic (FPGA can have 2 cores, then 10 cores)

Trapped between two ill defined areas - Hardware and Programming models

BOG 10 OS/R: list of key research challenges

Challenge X.1

Challenge X.2

BOG 10 OS/R Possible Research Directions

Summary

GRAPHIC?

PRD X.1 - direction 1

PRD X.2 - direction 2

PRD X.n - direction n

PRD X.n : Short title of possible research direction

- One paragraph description (3 sentence/bullet)
- Research challenges
 - Metrics for progress
- Potential research approaches and research directions
- How and when will success impact technology?